

# 9 Inference Without the Taking Condition

*Declan Smithies*

## 9.1. The Taking Condition

What is involved in making an inference? This chapter argues against what Paul Boghossian calls *the Taking Condition*:

**The Taking Condition:** Inferring necessarily involves the thinker *taking* his premises to support his conclusion and drawing his conclusion *because* of that fact.

(2014: 5)

I won't argue that the Taking Condition is *incoherent*—that nothing can coherently play the role that takings are supposed to play in inference. Instead, I'll argue that it cannot plausibly explain all the inferential knowledge that we ordinarily take ourselves to have. Moreover, I'll argue that we don't need it to understand the nature of inference.

It's worth noting from the outset that the Taking Condition doesn't explain the nature of inference in more basic terms. Instead, it presupposes the concept of *drawing a conclusion* without any further analysis. Moreover, it doesn't explain what it takes to draw a conclusion *because* you take it to be supported by your premises. Even so, the Taking Condition imposes a substantial and disputable requirement on making an inference; namely, that you must *take* your premises to support your conclusion. What exactly does this mean?

We cannot interpret this as a mere placeholder for whatever relation holds between your attitudes toward the premises and the conclusion when you make an inference. Otherwise, taking some premises to support a conclusion means nothing other than inferring the conclusion from those premises. I will ignore this deflationary interpretation since it is entirely vacuous. Instead, my goal is to argue against a representational interpretation of the Taking Condition, which says that the thinker—and not just her sub-personal systems—must *represent* that the premises of an inference support its conclusion.

On the view I defend, you can infer a conclusion from some premises without representing that the premises support the conclusion. Perhaps you must *treat* your premises as if they support your conclusion, but only in the trivial sense that you infer the conclusion from the premises. What I deny is that your inference from premises to conclusion must be mediated by some personal-level representation that those premises support the conclusion.

## 9.2. Regress Problems

What kind of representational state is involved in taking your premises to support a conclusion? On a *doxastic* construal of the Taking Condition, making an inference requires drawing a conclusion because you *believe* it is supported by your premises. As Boghossian argues, however, the doxastic construal generates what he calls an *ingress regress*. If inferential knowledge is possible, then some inferences can be justified. And yet the belief that some premises support a conclusion cannot play any role in justifying inference unless the belief is justified. So how is it justified? If it can be justified only by inference, then we face an infinite regress: you cannot be justified in making an inference unless you're antecedently justified in making an infinite series of inferences beforehand. And this makes it impossible for inference to be justified at all.

To avoid the ingress regress, Boghossian (2018: 62) endorses an *intuitional* construal of the Taking Condition, according to which inference involves drawing a conclusion from some premises because you intuit that the premises support the conclusion. An *intuition* is a conscious representational state—a seeming, appearance, or presentation—that is distinct from belief and that can justify belief without standing in need of any justification. Because intuition—unlike belief—needs no justification, the ingress regress never gets started.

In earlier work, Boghossian (2014: 9) argues that the intuitional view faces a regress problem of its own, which he calls the *egress regress*. The problem is that any representation that some premises support a conclusion—whether doxastic or not—must be combined with a representation of the premises before it can be used in inferring the conclusion from those premises. But this transition has the form of an inference:

- (1) P supports Q.
- (2) P.
- (3) Therefore, Q.

Hence, inference seems required to bring the intuition that some premises support a conclusion to bear in the act of inferring the conclusion from

those premises. Once again, this generates an infinite regress in which making an inference requires having already made an infinite series of inferences beforehand.

More recently, however, Boghossian suggests that we can block the egress regress by appealing to mundane examples in which representational states with conditional contents are used in rationally guiding action without any need for inference:

For example: the tennis player's intention to respond with a backhand if the ball comes to his left, and with a forehand if it comes to his right. That general intention can rationally control the tennis player's behavior without his having to conduct personal-level inferences in order to act upon it.

(2018: 63)

However, Boghossian's use of this example is questionable. When the tennis player prepares to hit a backhand, he does so because he *believes* that the ball is coming to his left. If he didn't believe this, then he wouldn't have formed the intention to hit a backhand. His unconditional belief that the ball is coming to his left combines with his conditional intention to hit a backhand when the ball comes to his left to generate the unconditional intention to hit a backhand now. And this causal transition has the form of an inference:

- (1) If the ball comes to my left, then I'll hit a backhand.
- (2) The ball is now coming to my left.
- (3) So, I'll hit a backhand now.

Moreover, this inferential transition operates on personal-level representations—namely, beliefs and intentions—rather than representations in some sub-personal system. So, pace Boghossian, the tennis player needs to make a personal-level inference in acting on his conditional intention. Otherwise, we cannot explain why his action is rational in the context of the game.

Of course, he cannot make the inference *consciously* by thinking through the premises in sequence: that would take too much time. As I argue in Section 9.3, however, not all personal-level inferences are made consciously. Inference requires only the right kind of causal dependence between personal-level representations of premises and conclusion. Perhaps a representational state is properly attributed to the person, rather than one of his sub-personal systems, only if its contents are accessible to consciousness (Smithies 2019: ch. 4). Even so, it doesn't follow that their contents must be consciously accessed in the act of making an inference.

Perhaps Boghossian will deny that the tennis player draws an inference because he violates the Taking Condition: he doesn't take his premises to support his conclusion. And yet it remains to be seen whether this response can be motivated on independent grounds. I have not given any principled argument that it's incoherent to posit personal-level representational states with conditional contents that can be used in non-inferentially guiding inference. Even so, we've not yet seen any plausible precedent for this controversial idea. If we can, we should try to make sense of inference without it.

### 9.3. Over-intellectualization Problems

Another problem with the Taking Condition is that it imposes overly demanding intellectual requirements on making an inference. As a result, it struggles to accommodate the possibility of unreflective reasoning in non-human animals, human children, and even human adults.

Adult human life involves participation in a social practice that is sometimes called "the game of giving and asking for reasons." We challenge each other to articulate reasons for our beliefs and actions, and we evaluate whether those reasons are any good. In responding to such challenges, we give arguments in which we cite considerations that we take to support our beliefs and actions. In evaluating these arguments, we explicitly consider whether the cited premises support their conclusions. This is what Tyler Burge (1996) calls *critical reasoning*.

As Burge notes, however, not all reasoning is critical reasoning. We often make inferences without evaluating them first. Such inferences can extend our knowledge without any need to represent the support relation between premises and conclusion. As I'll explain, these unreflective inferences can be fully conscious, partially conscious, or fully unconscious.

First, some unreflective reasoning is *fully conscious* in the sense that we consciously represent the premises and conclusion of an argument, although we don't consciously represent the support relation between them. Suppose you learn that the weather forecast predicts rain tomorrow and this prompts you to wonder whether your friend will cancel the bike ride that you planned together. Conscious thought might be needed to figure out how your friend will react to the weather forecast. But once you decide that they will cancel if it rains, and you already know that it will rain, you don't need to consider whether these premises support the conclusion that the ride will be cancelled. The conclusion will be evident without any further consideration at all.

Second, some reasoning is *partially conscious* in the sense that you consciously represent some premises of an argument but not all of them. If you already know from previous experience that your friend will cancel the

bike ride if it rains, then no further conscious thought is needed to figure out how they will react to the weather forecast. You can deduce that they will cancel the ride without first needing to recall your background knowledge that they will cancel if it rains. You can draw on this background knowledge in making the inference without accessing its content in conscious thought. Indeed, there are principled reasons why you cannot bring all your background knowledge into consciousness in the act of making an inference. Attention limits the capacity of conscious thought: you cannot think consciously about too many different things at the same time. But the background knowledge that is relevant to making an inference can overflow these attentional limits on conscious thought.

Third, some reasoning is *fully unconscious* in the sense that the conclusion is drawn from the premises without any conscious representation at all. You might go to bed wondering how your friend will react to the weather forecast and wake up already knowing that they will cancel. In that case, your knowledge is based on unconscious inference from background knowledge about your friend. This background knowledge may be consciously accessible in the sense that you can access it on demand when you wake up and ask yourself how you know they will cancel. But none of this background knowledge needs to be consciously accessed in the process of drawing the inference. If reasoning can be partially conscious, then it can be wholly unconscious for the same reason—your knowledge can figure in inference without conscious access.

If Boghossian denies that these are genuine inferences, then how can he explain your knowledge that the ride will be cancelled? One option is to deny that you have *knowledge*, rather than mere reliable belief (Boghossian 2019: 122). But why make this concession to skepticism? Boghossian sometimes invokes an internalist condition for knowledge, which requires that your reasons for belief must be in principle knowable by reflection alone:

On internalist ways of thinking . . . you have most reason to believe something only if it is possible for you to figure out, by reflection alone, that you have reason to believe it. And that in turn requires that, in the ideal case, all the factors that are relevant to assessing your reasoning should be open to reflective view.

(2016: 49)

And yet a weak version of the internalist requirement is satisfied so long as your conclusion is inferred from known premises that are accessible to consciousness upon demand. You don't need to actually reflect on and consciously access your reasons for belief in the act of making an inference. This stronger version of the internalist requirement is too demanding

because it excludes the possibility of unreflective knowledge altogether (Smithies 2019: 264–268).

A second option is to avoid this concession to skepticism by explaining your knowledge in some other way. Thus, Boghossian suggests that perhaps your knowledge is based on intuition, rather than inference:

The most natural description of the case where  $p$  suddenly strikes you as true is that you suddenly have the intuition that  $p$  is true. No doubt there is a causal explanation for why you suddenly have that intuition. And no doubt that causal explanation has something to do with your prior experiences with  $p$ . But all of that is a far cry from saying that you *inferred* to  $p$  from premises that you are not aware of.

(2019: 122)

This suggestion fails to explain how your knowledge that the ride will be cancelled depends on your knowledge about the weather forecast together with your knowledge of how your friend will react. Boghossian claims that the dependence is causal, rather than epistemic, because intuitions are caused rather than justified by background beliefs. What this ignores, however, is that your background beliefs must constitute *knowledge* to serve as a causal basis for knowing that the ride will be cancelled. Hence, your knowledge that the ride will be cancelled depends epistemically, and not just causally, on your background knowledge. We cannot plausibly explain this while denying that it's a genuine case of inference.

A third option is to accommodate these examples of inferential knowledge by positing *tacit* representations to satisfy the Taking Condition. Boghossian (2018: 66) claims that unreflective reasoning is guided by some tacit representation of support relations that is not consciously accessed in the act of making the inference. But what are these tacit representations? They cannot be sub-personal representations, since the Taking Condition requires that support relations are represented by thinkers themselves, rather than their sub-personal systems. And they are not conscious experiences either. So, presumably, they are dispositional beliefs.

I'm willing to concede that reflective adults typically have such dispositional beliefs when they make inferences. After all, it is irrational to make an inference while disbelieving or suspending belief about whether the premises support the conclusion: this kind of Moorean incoherence is an instance of epistemic akrasia (Hlobil 2014; cf. Smithies 2019: ch. 9). While not strictly impossible, this extreme form of irrationality is rather unusual. We're normally disposed to defend our beliefs in response to challenges by citing the premises from which we infer them. In other words, we tend to believe at least in a dispositional sense that the premises of our inferences support their conclusions.

Even so, the Taking Condition requires that our reasoning must be *guided* by these dispositional beliefs about the support relation. This is what happens in critical reasoning. In some cases, however, the direction of explanation is the other way around. For instance, we sometimes use suppositional reasoning as a strategy for evaluating arguments: we assume that their premises are true, and we evaluate the conclusion in light of these assumptions. In such cases, we may believe that some premises support a conclusion only because we're disposed to infer the conclusion from the premises in the context of suppositional reasoning.

Moreover, there are independent reasons to doubt that the Taking Condition can be salvaged by invoking dispositional beliefs. As we've seen, the doxastic construal of the Taking Condition generates an infinite regress of justification. To block the regress, Boghossian appeals to intuitions that can play a role in justifying inference without needing justification themselves. Intuitions are conscious experiences with the phenomenal character of seemings, appearances, or presentations. Dispositional beliefs have no such phenomenal character. Perhaps they are associated with dispositions to experience the phenomenal character of intuition. And yet the mere disposition to experience an intuition—unlike the experience of having one—plays no role in justifying belief.

An additional problem is that you need the *concept* of support to have a belief or intuition that some premises support a conclusion. Plausibly, however, unreflective creatures can make inferences without any conceptual capacity to represent support relations. Consider a child who infers that her father is home when she hears the front door open. Presumably, she needs no logical or epistemic concepts—such as deductive entailment or probabilistic support—to make the inference. Boghossian (2018: 67) speculates that children may acquire a generic concept of support expressible by the word “so” around 3–4 years old as part of coming to understand the distinction between appearance and reality. And yet toddlers with no more than a few simple words can make the inference in question when they respond to the door opening by saying, “Dada!” And even non-human animals, such as domestic dogs, can display inferential knowledge of who is at the door through various forms of non-linguistic behavior, such as grabbing a ball to play catch (see Andrews 2020: ch. 4).

Boghossian (2018: 61) seems forced to deny that these are genuine cases of inference. But it's implausible to deny that toddlers or pets can know who is at the door. After all, it's extremely natural to ascribe this knowledge to explain why they act as they do. And it's not clear how else they can acquire this knowledge except by inference. To deny that they have inferential knowledge is to make an implausible concession to skepticism. So Boghossian needs strong reasons to deny that animals and children can acquire knowledge through inference.

His main argument is that you can be held *responsible* for your reasoning. When you make an inference, your belief in the conclusion is based on the reasons provided by the premises from which you inferred it. We can evaluate the quality of the inference depending on whether it is based on good or bad reasons. Moreover, we hold people responsible for the quality of their reasoning: we regard good reasoning as praiseworthy and bad reasoning as blameworthy. Interestingly, however, we don't hold animals or toddlers responsible for their beliefs. Does this show that they don't engage in inference at all?

No: the argument proves too much. After all, animals and children are agents who act for reasons, whether good or bad. Even so, we don't hold them responsible for their actions by subjecting them to reactive attitudes, such as praise and blame. Arguably, responsibility requires the capacity for reflection (Smithies 2019: 280–282). You cannot be held responsible for your beliefs and actions unless you're capable of regulating them in light of reflection on your reasons for belief and action. We don't hold animals and toddlers responsible for their beliefs and actions because they lack these reflective capacities. But that doesn't show that they cannot believe and act for good or bad reasons at all. In particular, it doesn't show that they cannot make inferences by believing things for reasons provided by their other beliefs.

I don't claim to know where inference first makes its appearance on the phylogenetic or ontogenetic chain. That is an empirical issue that goes beyond the scope of this chapter. Even so, it seems overwhelmingly likely that this occurs before the capacity to represent support relations. Presumably, you cannot have the concept of support without the capacity to think thoughts of the form, "*p* supports *q*." But attention imposes limits on the capacity of thought: there are only so many thoughts we can entertain at once. For some creatures, these attentional limits may be severe enough that they can only think one atomic thought at a time. I see no principled reason to deny that such creatures can move inferentially from one thought to another without representing support relations between them. In any case, the burden is on proponents of the Taking Condition to explain why this is impossible.

#### 9.4. Generalization Problems

As we've just seen, our actions are based on reasons in much the same way as our beliefs. Just as we believe things for reasons—whether good or bad—so too we act and react for reasons. In other words, the basing relation extends beyond the epistemic domain of belief and knowledge into the practical domain of action and reaction.

Inference is just one species of the basing relation. When you make an inference, your belief in the conclusion is based on the reason provided by



the premises from which you infer it. But not all beliefs are inferentially based on other premises that you believe. Otherwise, we face an infinite regress in which every justified belief is inferentially based on reasons provided by some other justified belief and so on without end. Where does the regress come to an end? The usual foundationalist answer is that justified beliefs can be non-inferentially based on reasons provided by perception, introspection, or intuition. On this view, beliefs can be based on reasons either with or without inference.

What this means is that an account of inference must cohere with a more general account of the basing relation. After all, inference is just one species of the genus. So, if we endorse the Taking Condition, then we face a choice: either we must extend it to other species of basing or we must give special reasons for restricting it to inference. Unfortunately, however, neither option seems attractive.

We cannot plausibly extend the Taking Condition to all instances of the basing relation. We sometimes act, react, and believe things for reasons without representing them *as* reasons that support our response. Here are some examples:

- You can choose something (e.g., chocolate ice-cream) for the reason that you like it without representing that the fact that you like it is a reason to choose it.
- You can feel sad for the reason that you've lost something you value (e.g., a toy) without representing that the loss is a reason to feel sad.
- You can believe something (e.g., that the sky is blue) for the reason that it looks that way without representing that the visual appearance is a reason to believe it.

Young children—and non-human animals too—can make choices, feel emotions, and believe things for reasons without representing their reasons *as* reasons for responding as they do. Even mature adults often find themselves in this predicament. Although we have the conceptual capacities required for articulating our reasons for belief and action in general, we're not always capable of exercising these capacities accurately on any given occasion. Sometimes, we act for one reason while believing we act for another. As we know from social psychology, examples of self-deception, confabulation, and post-hoc rationalization are all too common (Wilson 2004).

Boghossian (2019: 122–123) resists extending the Taking Condition to all instances of the basing relation. He allows that you can acquire perceptual knowledge based on perceptual appearance without taking the appearance to support the belief. But why should the requirements for inferential basing diverge from the requirements for non-inferential basing? According to Boghossian, the Taking Condition doesn't apply when

there is a match in content between the justifying state and the belief that it justifies. He continues:

The inference from  $p$  to  $q$  is not like that. A belief that  $p$ , which is the input into the inferential process, is not a seeming that  $q$ . And while the transition to believing that  $q$  may be familiar and well-supported, it is not simply like acquiescing in something that is already the proto-belief that  $q$ .

(2019: 123)

While there is indeed a structural difference between inferential and non-inferential justification, I don't see how this motivates any restriction on the Taking Condition. The Taking Condition constrains inference whether justified or unjustified: to make an inference, you must take your premises to support believing your conclusion. So why doesn't the non-inferential transition from perception to belief require taking your perception to support your belief? This question isn't settled by noting the structural difference between inferential and non-inferential justification. We need some other reason to restrict the Taking Condition.

### 9.5. Reasoning as Mental Action

A more promising answer is that the Taking Condition applies only to instances of the basing condition that involve *agency*. Agency is not typically involved in forming beliefs that endorse the contents of perception: this is something we do automatically and involuntarily. In contrast, Boghossian maintains that reasoning is a form of mental action:

Reasoning is something we *do*, not just something that happens to us. And it is something *we* do, not just something that is done by sub-personal bits of us. And it is something that we do with an *aim*—that of figuring out what follows or is supported by other things one believes. It's hard to see how to respect these features of reasoning without something like the Taking Condition.

(2014: 5)

Moreover, the Taking Condition is supposed to explain *why* reasoning is a form of mental action. The general idea is that all actions are based on some desire, goal, or aim to achieve an end together with some belief-like representation that the action is a means to this end. Here is Boghossian's succinct summary of this idea:

The agent has an aim; she has a view about a way of accomplishing that aim; and she performs an action as a result of that combination.

(2018: 61–62)

Reasoning fits this mold when it is based on the aim of figuring out what follows from some premises together with a representation that inferring a given conclusion will achieve the aim because it follows from those premises. In that case, the Taking Condition is satisfied.

Jonathan Way and Daniel Whiting (2016: 325) complain that Boghossian's explanation is viciously regressive. According to Boghossian, every action is based on the aim of achieving some end together with a representation that the action is a means to this end. But this seems tantamount to saying that every action is based on practical reasoning of the following form:

- (1) I shall do A.
- (2) Doing B is a means to doing A.
- (3) So, I shall do B.

On Boghossian's view, however, all reasoning—including practical reasoning—is mental action. This yields the problematic result that every action requires some prior mental action, and all reasoning requires some prior reasoning. Something has clearly gone awry.

To avoid the regress, Boghossian must deny that all action is based on practical reasoning. This seems like a sensible thing to deny. Presumably, a toddler can choose chocolate ice-cream for the reason that they prefer it without reasoning from the premise that they prefer it together with the further premise that choosing it will satisfy their preference. If he denies this, however, then he needs to explain how an action can be non-inferentially based on an aim, goal, or end. We already encountered this challenge in Section 9.2.

Even if this challenge can be met, however, another regress problem looms nearby. According to Boghossian, all action is based—whether inferentially or otherwise—on some end together with some representation of the means to that end. Now the question arises whether all forms of basing involve agency. If so, then we face an infinite regress: every action is based on some prior act of basing. Hence, Boghossian must deny that all basing involves agency. But if agency is not required for basing in general, then why should it be required for *inferential* basing in particular? Ultimately, I'm not persuaded that Boghossian provides any compelling answer to this question.

His opening gambit is that reasoning is something that we can be said to do, rather than something that happens to us. As a linguistic point, this seems entirely correct. At the same time, however, it cannot bear much theoretical weight. We often describe events that are not actions in the active voice as things we do, rather than things that merely happen to us. For instance, we describe people as falling asleep, sweating, and vomiting,

as well as perceiving things and believing that things are thus-and-so. Although these are all things we can be said to do, the Taking Condition does not apply to them. So, we cannot rely on this linguistic point in making the case that inference is distinguished from non-inferential basing by its agential character.

A more promising consideration is that reasoning involves the *experience* of agency. When you concentrate on solving a puzzle, for example, your experience has an agential dimension: it feels as if you're actively trying to figure out the solution. As we saw in Section 9.3, however, not all reasoning is like this. Much of the time, we draw inferences effortlessly without any need to actively focus our attention. In such cases, there is no salient experience of agency. Our reasoning is often guided by background beliefs that are not manifest in consciousness at all. And we sometimes update our beliefs through reasoning that is wholly unconscious—as when we wake up from a deep sleep to find that we have already figured out the solution to a problem. Reasoning need not involve any experience at all, let alone an experience of agency.

Reasoning is one of many psychological phenomena that can be experienced either actively or passively. One such example is *imagination*: you can actively try to visualize a scene, or you can simply find yourself experiencing visual imagery in a dream or a daydream. Another is *attention*: you can focus your attention on following a moving target, or your attention can be captured when a stationary target suddenly moves. A third example is *judgment*: you can reach a verdict by actively assessing the evidence for and against a hypothesis, as in jury deliberations, but you can also experience a sudden realization as if out of the blue without any experience of agency. If all these other psychological processes can be experienced both actively and passively, then why not reasoning too?

I contend that reasoning is no different from perceptual belief-formation in this respect. We typically acquire perceptual knowledge automatically with no need to actively focus our attention on the question of what to believe. In some cases, however, we actively deliberate about whether to believe that things are how they perceptually appear. Similarly, our reasoning can be actively guided by focusing attention on whether one thing is a reason to believe another. But much of our reasoning occurs automatically without any need for active control.

Boghossian's main argument is that we cannot make sense of our practice of holding people responsible for their reasoning unless we suppose that reasoning is a mental action:

For it to make sense to hold you responsible for your inferences, inferring has to be something you do, and not just something that happens to you. It has to be a mental action of yours, something you have control

over, and which you could have done differently, had you thought it desirable to do so.

(2018: 60)

As Mark Richard (2019) notes, however, we hold people responsible not only for actions over which they have direct control but also for the consequences of their actions over which they have only indirect control. We can blame someone for ill health that results from bad decisions, for example, and we can praise them when they turn things around. So, we don't need to assume that reasoning is a mental action to explain why we hold people responsible. We just need to assume that people have indirect control over their reasoning by performing other mental actions, such as considering questions, paying attention, or gathering evidence.

This objection doesn't strike at the heart of Boghossian's argument. His main point is that we hold people responsible for the *reasons* on which their inferences are based. Although we hold people responsible for their health, we don't hold them responsible in the same way, since the state of your health is not based on reasons at all. A more charitable version of Boghossian's argument is that we need to assume that inference is a mental action to explain the distinctive way in which we hold people responsible for basing their inferences on good reasons.

The real problem with Boghossian's argument is that it proves too much. After all, we hold each other responsible for basing our beliefs on good reasons even when they are not based on inference. For instance, we credit people for basing their beliefs on perceptual appearances and we blame them when they fail to do so without good reason. If his argument succeeds at all, then it shows that all basing—whether inferential or non-inferential—is agential. As we've seen, however, this claim generates a vicious regress. I conclude that Boghossian cannot solve the generalization problem by appealing to the idea that inference is a mental action.

## 9.6. Inference and Association

One of Boghossian's main arguments for the Taking Condition is that we need it to explain the distinction between inference and association. Suppose you believe one thing because you believe another, although you don't take the one thing to support the other. What distinguishes this transition from a mere process of association that is sensitive to the contents of your beliefs? Here is Boghossian's example:

A habitual depressive's judging "I am having so much fun" may routinely cause and explain his judging "Yet there is so much suffering in

the world,” as directly as you please, without this being a case in which he is inferring the latter thought from the earlier one.

(2014: 4)

The depressive doesn't make an inference because his judgment that the world is full of suffering isn't based on his judgment that he is having fun. Rather, one judgment causes the other by a process of associative thinking. But then what more is needed for inference beyond the causation of one judgment by another?

The premises of an inference need not support its conclusion, since inference is not always justified. Instead, Boghossian claims, the thinker must *take* their premises to support their conclusion. The depressive doesn't draw an inference because he doesn't take his judgment that he is having fun to support his judgment that the world is full of suffering. In this way, Boghossian uses the Taking Condition to explain the distinction between inference and association.

Boghossian's example is somewhat under-described. On the most natural interpretation, the depressive already believes that the world is full of suffering, since his depression darkens his outlook on the world. The realization that he is having fun today merely serves as a trigger that activates this standing belief in conscious judgment. The same judgment might have been activated in some other way, such as considering whether the world is full of suffering. And this is enough to explain why the transition doesn't count as an inference. The one judgment merely activates the other, rather than serving as its epistemic basis. Hence, we can explain why the transition doesn't count as an inference without endorsing the Taking Condition.

Of course, Boghossian might deny that this version of the case is the one he has in mind. He could stipulate that the realization that you're having fun doesn't merely trigger the activation of a pre-existing belief, but rather causes the formation of a new belief. This version of the case is not so realistic—it has no obvious connection with the psychology of depression—but it is perfectly coherent all the same. The problem is that the example is not fleshed out in enough detail to make it clear why this should count as association, rather than inference. To be sure, the inference in question would be bizarrely irrational. And yet Boghossian imposes no limits on how irrational one's inferences can be. Rather than assuming from the outset that there are rationality constraints on inference, he requires that any such constraints should emerge as consequences of a correct theory of inference (Boghossian 2014: 4).

I'm not disputing that we can construct examples in which the causal transition from one belief to another is the wrong kind to constitute

inference. Indeed, examples of this kind are familiar from the literature on deviant causal chains. Here is one from John Turri:

Through some random quirk—the result of a neural assembly malfunctioning—Wilt’s belief that the lettuce has wilted is the proximate mental cause of his belief that the Patriots will win twelve games this season. But it certainly seems false that Wilt’s belief that the lettuce has wilted is his reason for believing that the Patriots will win twelve games this season.

(2011: 389)

Such examples show that inference is not just believing one thing because you believe another. To count as an inference, one belief must cause another in the right kind of way. The problem is that it’s hard to specify what counts as right kind of causal relation to constitute inference. This is the familiar problem of deviant causal chains. The problem of distinguishing between inference and association is just one aspect of this more general problem. Boghossian takes his challenge to be distinctive because it arises when the causal relation between beliefs is direct, rather than mediated by an intervening causal chain of events. As Turri’s example shows, however, mental causation can be deviant even when it is direct.

As Boghossian (2014: 5, n. 2) admits, the Taking Condition cannot solve the problem of deviant causal chains. It says that making an inference requires drawing a conclusion from some premises *because* you take the premises to support your conclusion, but it doesn’t specify which kind of causal relation is required for inference. Nothing precludes the representation of support relations from figuring in deviant causal chains. Suppose your belief in some conclusion is caused in some *deviant* way by your taking your premises to support the conclusion. This doesn’t count as an inference because your belief in the conclusion isn’t caused in the right way. And yet the Taking Condition doesn’t specify which kind of causal relation is required.

For much the same reason, the Taking Condition cannot explain the distinction between inference and association. If beliefs can figure in processes of content-sensitive association, then so can the representational state of taking some premises to support a conclusion. Suppose our depressive thinker mistakenly takes the premise that he is having fun to support the conclusion that the world is full of suffering. And suppose this causes him by a process of association to form some arbitrary belief—say, that he is the King of Spain. This is surely not an inference. If so, then the same is true when the same inputs cause him to believe by a process of association that the world is full of suffering. In this example, the Taking Condition is satisfied, since the thinker believes that Q because he believes that P and he

takes P to support Q. And yet this is a case of association, rather than inference. Hence, the Taking Condition cannot explain the distinction between inference and association after all.

To explain the distinction in full generality, we need to specify which kind of causal relation is necessary and sufficient for inference. In other words, we need a solution to the problem of deviant causal chains. Unfortunately, however, there is no consensus about how to solve this problem or even whether it can be solved at all. Perhaps we must take the concept of inference as primitive rather than try to analyze it in more basic terms.

As we've seen, Boghossian's Taking Condition presupposes the concept of drawing an inference *because* you take the premises to support the conclusion. He offers no analysis of what counts as the right kind of causal relation for inference. But if he can take this as primitive, then his opponents are entitled to do the same. The disputed issue is whether the relation of this causal relation must include representation of not only premises and conclusion but also the support relation between them. We've not yet seen any compelling motivation for this claim.

## 9.7. Conclusions

This chapter makes no attempt to explain what it is to make an inference in more basic terms. Indeed, I doubt that any reductive explanation of inference is possible. In any case, Boghossian's Taking Condition does no such thing: it merely states a requirement for making an inference. I've argued that this requirement is too demanding because it cannot explain all of the knowledge that we ordinarily take ourselves to have. Moreover, I've explained why Boghossian's main arguments for the Taking Condition are not persuasive. I conclude that the Taking Condition is false: you can make an inference without taking your premises to support your conclusion.

## References

- Andrews, K. 2020. *The Animal Mind: An Introduction to the Philosophy of Animal Cognition*. Routledge.
- Boghossian, P. 2014. What Is Inference? *Philosophical Studies* 169: 1–18.
- Boghossian, P. 2016. Reasoning and Reflection: A Reply to Kornblith. *Analysis* 76: 41–54.
- Boghossian, P. 2018. Delimiting the Boundaries of Inference. *Philosophical Issues* 28: 55–69.
- Boghossian, P. 2019. Inference, Agency, and Responsibility. In *Reasoning: New Essays on Theoretical and Practical Thinking*, edited by M. Balcerak Jackson & B. Balcerak Jackson, 101–124. Oxford University Press.
- Burge, T. 1996. Our Entitlement to Self-Knowledge. *Proceedings of the Aristotelian Society* 96: 91–116.



- Hlobil, U. 2014. Against Boghossian, Wright, and Broome on Inference. *Philosophical Studies* 167: 419–429.
- Richard, M. 2019. Is Reasoning a Form of Agency? In *Reasoning: New Essays on Theoretical and Practical Thinking*, edited by M. Balcerak Jackson & B. Balcerak Jackson, 91–100. Oxford University Press.
- Smithies, D. 2019. *The Epistemic Role of Consciousness*. Oxford University Press.
- Turri, J. 2011. Believing for a Reason. *Erkenntnis* 74: 383–397.
- Way, J. & Whiting, D. 2016. Against the Taking Condition. *Philosophical Issues* 26: 314–331.
- Wilson, T. 2004. *Strangers to Ourselves: Discovering the Adaptive Unconscious*. Harvard University Press.